User Manual - English

FUJITSU

InfiniBand Host Channel Adapter (100Gb) EDR

Linux

October 2015

Comments… Suggestions… Corrections…

The User Documentation Department would like to know your opinion of this manual. Your feedback helps us optimize our documentation to suit your individual needs.

Feel free to send us your comments by e-mail to manuals@ts.fujitsu.com.

Certified documentation according to DIN EN ISO 9001:2008

To ensure a consistently high quality standard and user-friendliness, this documentation was created to meet the regulations of a quality management system which complies with the requirements of the standard DIN EN ISO 9001:2008.

cognitas. Gesellschaft für Technik-Dokumentation mbH
www.cognitas.de

Copyright and Trademarks

Copyright © 2015 Fujitsu Technology Solutions GmbH.

Copyright © Fujitsu Technology Solutions GmbH 2015.

All rights reserved.
Delivery subject to availability; right of technical modifications reserved.

All hardware and software names used are trademarks of their respective manufacturers.

Contents

1 Introduction 5

1.1 Notational conventions 5
1.2 Target Group 6
1.3 Documentation Overview 6
1.4 Overview 6
1.5 Product Prerequisites 7

2 Important information 9

2.1 Safety instructions 9
2.2 ENERGY STAR 10
2.3 CE conformity 11

3 Adapter Card Hardware 13

3.1 Adapter Cards Covered in this Manual 13
3.2 Finding the GUID and serial number on the adapter cards . 15
3.3 Adapter Card Interfaces 16
3.3.1 I/O Interfaces 16
3.3.2 InfiniBand Interface 16
3.3.3 PCI Express Interface 17
3.3.4 LED Assignment 17
3.4 Adapter Card Installation 18
3.4.1 Hardware and Software Requirements 18
3.4.2 Bracket installation instructions 19
3.4.3 Installation instructions 20
3.4.4 Cables 20

4 Driver installation 23

4.1 Overview 23

4.2 Software and hardware requirements . . . 23
4.3 OFED installation . . . 24
4.3.1 Installing Mellanox OFED . . . 24
4.3.2 Port type management . . . 29
4.3.3 Uninstalling Mellanox OFED . . . 29
4.4 Updating Firmware after Installation . . . 30
4.5 Subnet Manager . . . 30
4.5.1 Functions of OpenSM . . . 30
4.5.2 Confirmation that OpenSM is Running . . . 31
4.5.3 Confirmation of the master SM . . . 31

5 Troubleshooting . . . 33

5.1 General . . . 33
5.2 Linux . . . 34

6 Appendix . . . 35

6.1 A.1 EMC Certification Statements . . . 35
6.1.1 FCC Statements (USA) . . . 35
6.1.2 EN Statements (Europe) . . . 36
6.1.3 ICES Statements (Canada) . . . 36
6.1.4 VCCI Statements (Japan) . . . 37
6.1.5 KCC Certification (Korea) . . . 37
6.2 Specifications . . . 38

1 Introduction

This manual describes the hardware, software and adapter card installation and basic use of the InfiniBand Host Channel Adapter (IB HCA).

HCA card name	Product name
PIB EP 100Gb 1 Port EDR	S26361-F5540-E102/-L102
PIB EP 100Gb 2 Port EDR	S26361-F5540-E202/-L202

Table 1: Product names

It also describes the installation of OFED (OpenFabrics Enterprise Distribution) for Linux.

1.1 Notational conventions

The following notational conventions are used in this manual:

Text in italics	indicates commands or menu items
`fixed font`	indicates system output
`semi-bold fixed font`	indicates text to be entered by the user
"Quotation marks"	indicate names of chapters and terms that are being emphasized.
►	describes activities that must be performed in the order shown.
Abc	indicates keys on the keyboard
CAUTION!	Pay particular attention to texts marked with this symbol! Failure to observe this warning may endanger your life, destroy the system or lead to the loss of data.
	indicates additional information, notes and tips.

Fujitsu Support

Please contact your Fujitsu Technical Support if you require assistance:

http://support.ts.fujitsu.com/Index.asp

If you purchased the products from Fujitsu Japan, please consult a system engineer.

1.2 Target Group

This manual is intended for users, developers, and system administrators responsible for setting up and maintaining switch system platforms using InfiniBand fabrics and networks. The switch system platforms must have a management server or client, in order for the package to work.

The manual assumes familiarity with the InfiniBand® Architecture Specification.

1.3 Documentation Overview

All documentation on PRIMERGY hardware and ServerView software is available online from the Fujitsu manuals server at:

- For the global market *http://manuals.ts.fujitsu.com*
- For the Japanese market: *http://jp.fujitsu.com/platform/server/primergy/manual*

The complete PRIMERGY documentation set can also be downloaded as a DVD ISO image at:

- For the global market: *ftp://ftp.ts.fujitsu.com/images/serverview/manuals*
- For the Japanese market: *http://jp.fujitsu.com/primergy/downloads*

1.4 Overview

This document is a user manual for InfiniBand Host Channel Adapters (IB HCA) card in rack servers. The card described in this manual has the following features:

- Single and Dual Port 100Gb/s IB HCA card

- Low Profile PCIe form factor with exchangeable Full Height (FH) or Low Profile (LP) bracket
- Based on the IB controller chip Mellanox ConnectX-4
- PCIe interface
  - PCIe x16 bus interface
  - PCIe 3.0 (8GT/s)
  - Auto-negotiates to x16, x8, x4, x2, or x1 (x8, x4, x2 or x1 will decrease the performance)
- IBTA Specification 1.3 compliant (plan)
- Hardware-based congestion control
- Connectivity
  - Interoperable with InfiniBand switches
  - QSFP connectors
  - 1m, 3m (100Gb/s) of passive copper cable

1.5 Product Prerequisites

System Class	Subsegment	Systems	Max.No.	Bracket LP/FH
All-round Server		**TX2560 M1**	**2**	**FH**
Rack Server	Entry	RX2530 M1	1	LP
		RX2540 M1	2	LP
	Midrange	RX4770M2	2	FH
Cloud Server		CX2550 M1	1	LP
		CX2570 M1	1	LP

LP = Low Profile bracket
FH = Full Height bracket

2 Important information

2.1 Safety instructions

The following safety instructions are also provided in the manual "Safety Notes and Regulations" or " 安全上のご注意 ".

This device meets the relevant safety regulations for IT equipment. If you have any questions about whether you can install the server in the intended environment, please contact your sales outlet or our customer service team.

Modules with Electrostatic-Sensitive Devices

Modules with electrostatic-sensitive devices are identified by the following sticker:

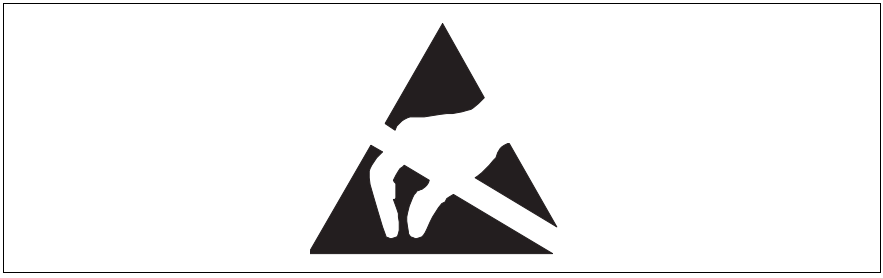

Figure 1: ESD label

When you handle components fitted with ESDs, you must always observe the following points:

- Switch off the system and remove the power plugs from the power outlets before installing or removing components with ESDs.
- You must always discharge static build-up (e.g. by touching a grounded object) before working with such components.
- Any devices or tools that are used must be free of electrostatic charge.
- Wear a suitable grounding cable that connects you to the external chassis of the system unit.
- Always hold components with ESDs at the edges or at the points marked green (touch points).

- Do not touch any connectors or conduction paths on an ESD.
- Place all the components on a pad which is free of electrostatic charge.

For a detailed description of how to handle ESD components, see the relevant European or international standards (EN 61340-5-1, ANSI/ESD S20.20).

Disposal

Disposal of this equipment should be in accordance to all national laws and regulations.

Other important information:

- Keep this operating manual and the other documentation (such as the Technical Manual, DVD) close to the device. All documentation must be included if the equipment is passed on to a third party.

2.2 ENERGY STAR

Products that have been certified compliant with ENERGY STAR and identified as such are in full compliance with the specification at shipping. Note that energy consumption can be affected by software that is installed or any changes that are made to the hardware configuration or BIOS or energy options subsequently. In such cases, the properties guaranteed by ENERGY STAR can no longer be assured.

The "ServerView Operations Manager" user guide contains instructions for reading out measurement values, including those relating to current energy consumption and air temperatures. Either the Performance Monitor or the Task Manager can be used to read out CPU utilization levels.

2.3 CE conformity

The system complies with the requirements of the EC directives 2004/108/EC regarding "Electromagnetic Compatibility" and 2006/95/EC "Low Voltage Directive" and the directive of the European Parliament and Council 2011/65/EU. This is indicated by the CE marking (CE = Communauté Européenne).

3 Adapter Card Hardware

3.1 Adapter Cards Covered in this Manual

The following table lists the adapter cards described in this manual:

HCA card name	PCI Express SERDES Speed	Data Transmission Rate/port	Bracket type	RoHS
PIB EP 100Gb 1 Port EDR	PCIe Gen3 8.0 GT/s	100Gb/s	Low profile and Full height	R-6
PIB EP 100Gb 2 Port EDR	PCIe Gen3 8.0 GT/s	100Gb/s	Low profile and Full height	R-6

Table 2: Single and Dual Port Card

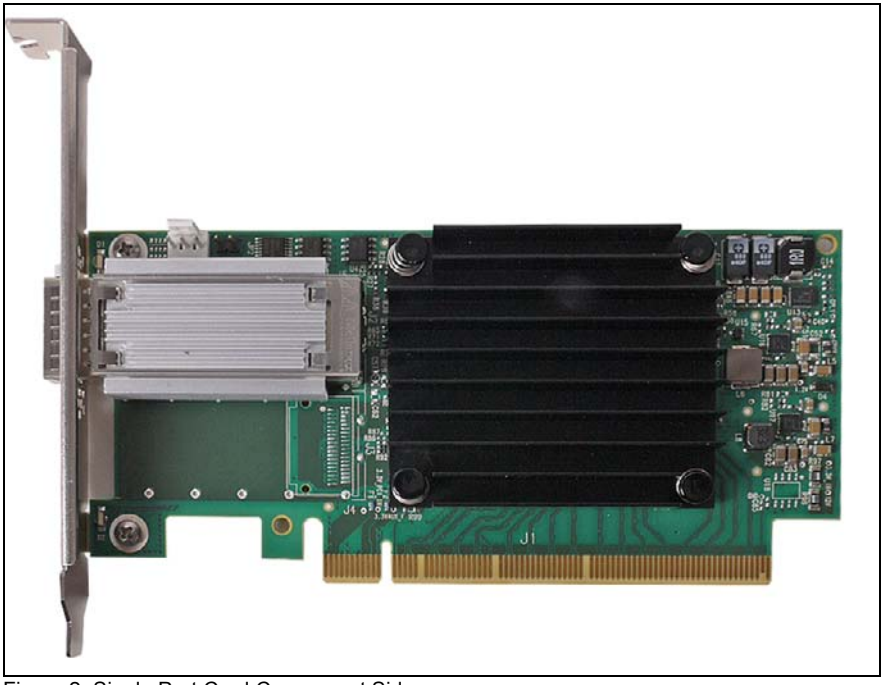

Figure 2: Single Port Card Component Side

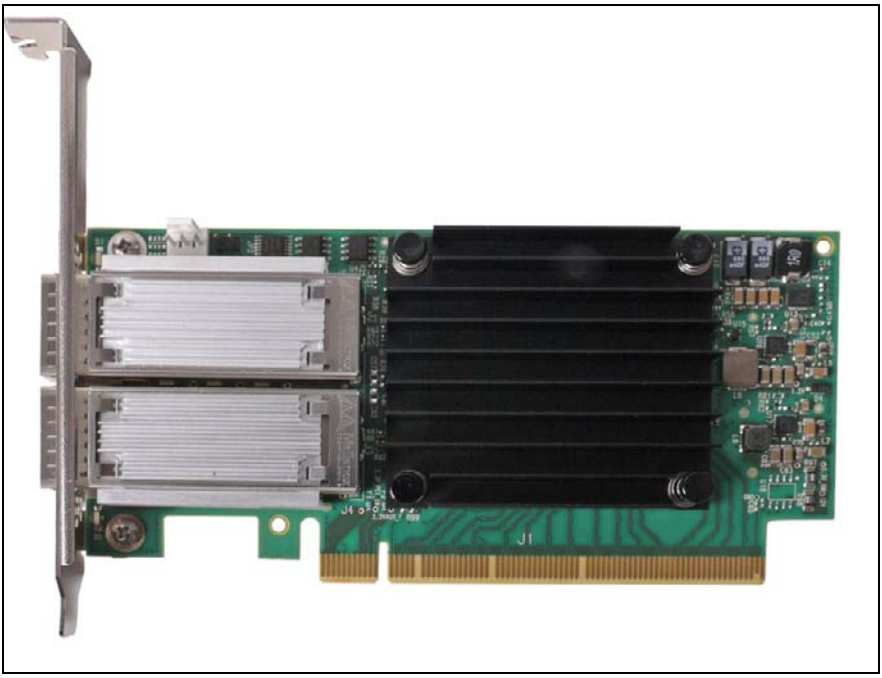

Figure 3: Dual port card component side

3.2 Finding the GUID and serial number on the adapter cards

All adapter cards have a label on the printed side of the adapter card that has the card serial number and the card GUID for InfiniBand protocol.

Figure 4: Card Product Label (e.g. IB HCA 1port EDR)

3.3 Adapter Card Interfaces

3.3.1 I/O Interfaces

Each adapter card includes the following interfaces:

– QSFP ports

– PCI Express x16 edge connector

– I/O panel LEDs

For dual port cards, port 1 connects to connector 1 of the device, while port 2 connects to connector 2 of the device.

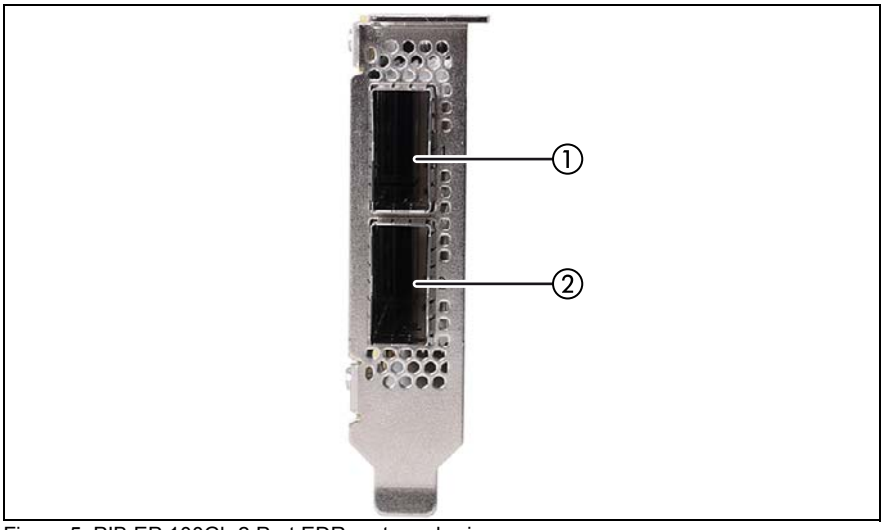

Figure 5: PIB EP 100Gb 2 Port EDR port numbering

3.3.2 InfiniBand Interface

The ConnectX-4 device is compliant with the InfiniBand Architecture Specification, Release 1.3.

Adapter cards (listed in table 2 on page 13) based on this device provide access to its ports by means of QSFP connectors.

3.3.3 PCI Express Interface

The ConnectX®-4 adapter cards support PCI Express 3.0 (1.1, 2.0 compatible) through an x16 edge connector. The device can be either a master initiating the PCI Express bus operations or a slave responding to PCI bus operations.

3.3.4 LED Assignment

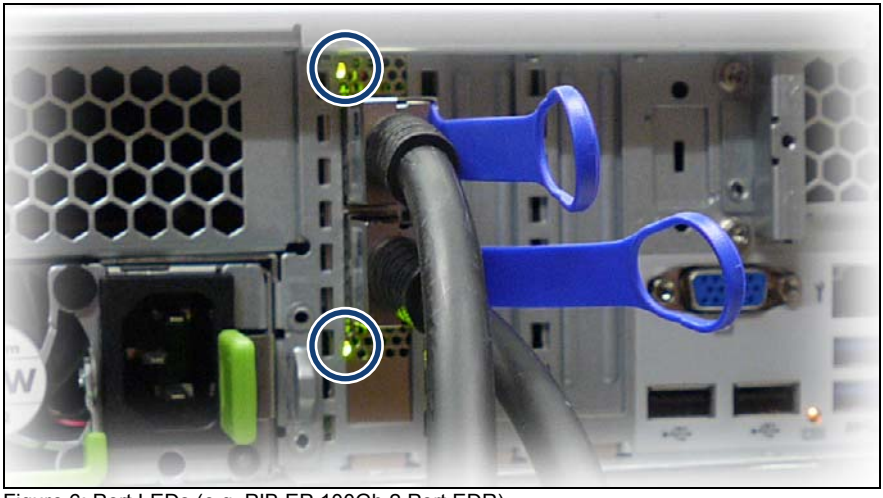

Figure 6: Port LEDs (e.g. PIB EP 100Gb 2 Port EDR)

There is one bi-color I/O LED per port (see circles).

LED	Function
Yellow - physical link	– Constant on indicates a good physical link – Blinking indicates a problem with the physical link – If neither LED is lit, then the physical link has not been established
Green - logical (data activity) link	– A blinking green indicates a valid logical link with data transfer

3.4 Adapter Card Installation

CAUTION!

Follow the safety instructions in the chapter "Important information" on page 9.

The adapter is being installed in a system that operates with voltages that can be lethal. Before opening the case of the system, observe the following precautions to avoid injury and prevent damage to system components.

- Remove any metallic objects from your hands and wrists.
- Make sure to use only insulated tools.
- Verify that the system is powered off and is unplugged.
- It is strongly recommended to use an ESD strap or other antistatic devices.

3.4.1 Hardware and Software Requirements

Before installing the IB HCA card, please make sure that the system meets the hardware and software requirements listed in table 3.

Requirement	Description
Hardware	– Minimum 3 GB of memory – PCI Express x16 slots
Software Operating Systems/Distributions	For Linux, the InfiniBand drivers are in the Open Fabrics Enterprise Distribution (OFED) software package which is available at *http://support.ts.fujitsu.com/download/* Drivers & Downloads If you purchased the products from Fujitsu Japan, you can download it from *http://jp.fujitsu.com/platform/server/primergy/downloads/*

Table 3: Hardware and software requirements

3.4.2 Bracket installation instructions

The card is usually shipped with a tall bracket installed. If this form factor is suitable for your requirements, you can skip the remainder of this section and move to section "Installation instructions" on page 20. If you need to replace it with the short bracket that is included in the shipping box, please follow the instructions in this section.

- Remove the mounting bracket screws from the top of the adapter (see circles in figure 7).

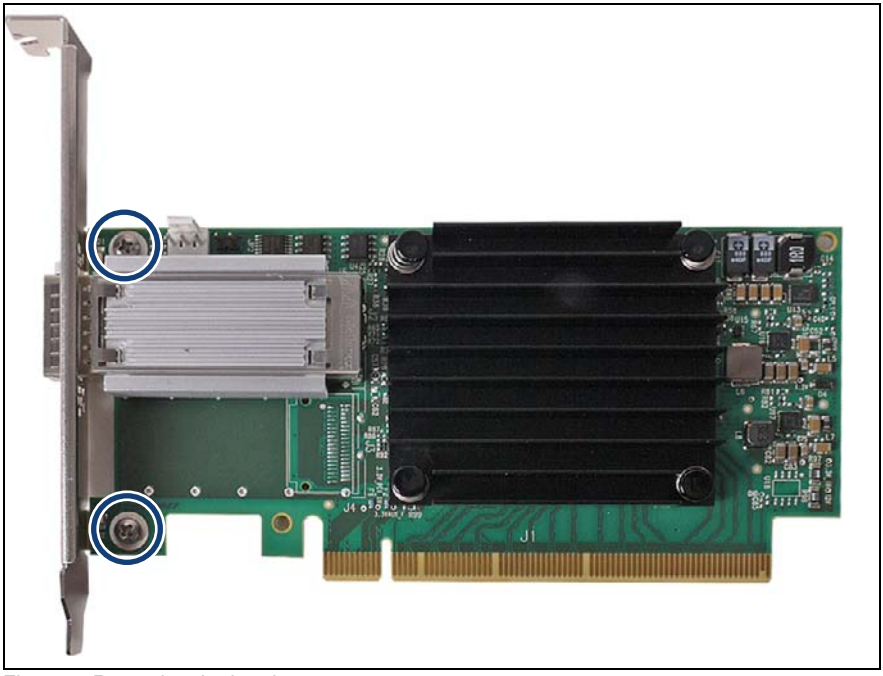

Figure 7: Removing the bracket

- Remove the two screws holding the bracket in place. The bracket comes loose from the card.
- Save the two screws with the captive washers.
- Align the new mounting bracket tabs with the holes in the adapter.

> **i** Do not force the bracket onto the card. Be careful not to damage the socket's grounding tabs.

► Re-install the screws that attach the adapter to the bracket.

The bracket of the dual port HCA is replaced in the same way.

3.4.3 Installation instructions

Read all installation instructions before connecting the equipment to the power source.

The adapter cards listed in table 2 on page 13 are standard PCI Express cards, each with a standard x16 edge connector. Please consult the host machine documentation for instructions on how to install a PCI Express card.

CAUTION!

When more than one PCI slot is available make sure you use the PCI slot with the proper configuration.

Any PCI slot with the proper configuration is acceptable for connection. If the card is installed in a PCI slot with less lanes than the card requires, then the adapter card will not provide optimum data transfer.

3.4.4 Cables

The QSFP port can be connected to InfiniBand passive copper and active optical cables.

3.4.4.1 Cable Installation

All cables can be inserted or removed with the unit powered on. To insert a cable, press the connector into the port receptacle until the connector is firmly seated. The LED indicator will light yellow when the physical connection is established (that is, when the unit is powered on and a cable is plugged into the port with the other end of the connector plugged into a functioning port). After plugging in a cable, lock the connector using the latching mechanism particular to the cable vendor. When a logical connection is made the LED indicator will light green. When data is being transferred, the LED indicator blinks green.

CAUTION!

When installing cables make sure that the latches engage.

CAUTION!

Always install and remove cables by pushing or pulling the cable and connector in a straight line with the card.

Care should be taken not to impede the air exhaust flow through the ventilation holes. Cable lengths should be used which allow for routing horizontally around to the side of the chassis before bending upward or downward in the rack.

To remove a cable, disengage the locks and slowly pull the connector away from the port receptacle. Both LED indicators will turn off when the cable is unseated.

CAUTION!

Cables, especially long copper cables, can weigh a substantial amount. Make sure that the weight of the cable is supported on its own and is not hanging from IB HCA card.

3.4.4.2 Inserting a Cable into the Adapter Card

- Support the weight of the cable before connecting the cable to the adapter card. Do this by using a cable holder or tying the cable to the rack.
- Determine the correct orientation of the connector to the card before inserting the connector. Do not try and insert the connector upside down. This may damage the adapter card.

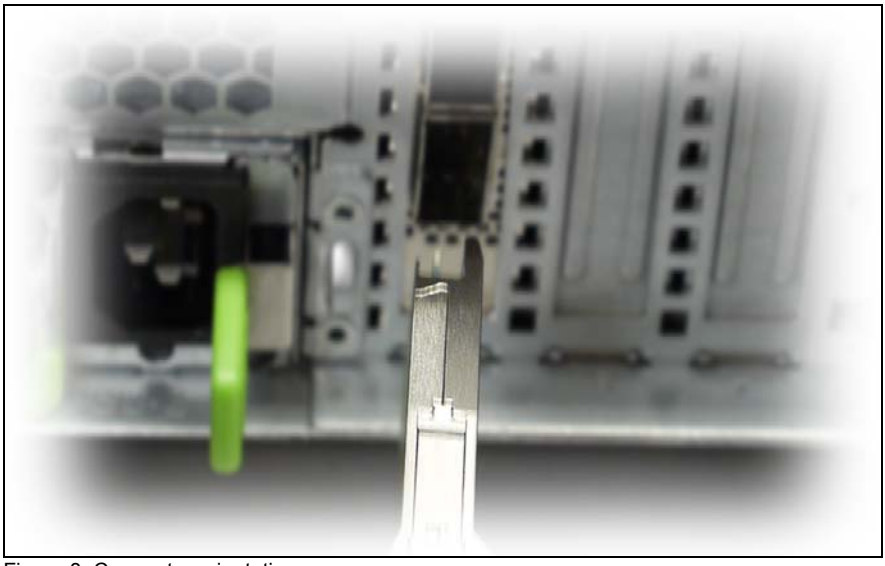

Figure 8: Connector orientation

- Insert the connector into the adapter card. Be careful to insert the connector straight into the cage. Do not apply any torque, up or down, to the connector cage in the adapter card.
- Make sure that the connector locks in place.

3.4.4.3 Removing a Cable from the Adapter Card

1. Pull on the latch release mechanism to unlatch the connector and pull the connector out of the cage.
2. Do not apply torque to the connector when removing it from the adapter card.
3. Remove any cable supports that were used to support the cable's weight.

4 Driver installation

4.1 Overview

This chapter describes how to install a single host machine with Fujitsu InfiniBand hardware installed. A server can be properly installed with all required InfiniBand drivers and software during Linux installation or afterwards by using the OpenFabrics Enterprise Distribution (OFED).
In the first case, the correct packages must be selected during setup and should only be performed by an advanced user. For simple complete installation, use the OFED drivers.

If you purchased the products from Fujitsu Japan, you can download it from *http://primeserver.fujitsu.com/primergy/downloads/*.

The chapter includes the following sections:

- section "OFED installation" on page 24
- section "Updating Firmware after Installation" on page 30
- section "Subnet Manager" on page 30

4.2 Software and hardware requirements

Required Disk Space for Linux Installation

1 GB

PCI Device ID

– 15B3:1013:15B3:0030

– 15B3:1013:15B3:0032

Operating system

– RHEL (Red Hat Enterprise Linux) 6.6, 6.7, 7.1 or higher

– SLES11 (SuSE Suse Linux Enterprise Server) SP3, SP4, SLES12 or higher

Installer privileges

The installation requires administrator privileges on the target machine.

4.3 OFED installation

Please install OFED by the following steps after the OS installation by SVIM (Server Installation Manager).

- Download OFED from Fujitsu Technology Solutions Web site *http://support.ts.fujitsu.com/*
- If you purchased the products from Fujitsu Japan, you download it from *http://jp.fujitsu.com/platform/server/primergy/downloads/*.

CAUTION!

This software is the driver package of ConnectX-4 IB HCA card (100Gb).

This package driver does not support ConnectX-3 IB HCA card (56Gb).

When OFED downloaded from the above is being used, please don't apply the RHEL native OFED at the time of RHEL update.

4.3.1 Installing Mellanox OFED

The installation script, mlnxofedinstall, performs the following:

- Discovers the currently installed kernel
- Uninstalls any software stacks that are part of the standard operating system distribution or another vendor's commercial stack
- Installs the MLNX_OFED_LINUX binary RPMs (if they are available for the current kernel)
- Identifies the currently installed InfiniBand and Ethernet network adapters and automatically upgrades the firmware

Usage

`./mnt/mlnxofedinstall` *OPTIONS*

The installation script removes all previously installed Mellanox OFED packages and re-installs from scratch. You will be prompted to acknowledge the deletion of the old packages.

Pre-existing configuration files will be saved with the extension *.conf.rpmsave.*

– If you need to install Mellanox OFED on an entire (homogeneous) cluster, a common strategy is to mount the ISO image on one of the cluster nodes and then copy it to a shared file system such as *NFS*. To install on all the cluster nodes, use cluster-aware tools (such as *pdsh*).

– If your kernel version does not match with any of the offered pre-built RPMs, you can add your kernel version by using the *mlnx_add_kernel_support.sh* script located under the *docs/ directory*.

On Redhat and SLES distributions with errata kernel installed there is no need to use the *mlnx_add_kernel_support.sh* script. The regular installation can be performed and weak-updates mechanism will create symbolic links to the *MLNX_OFED* kernel modules.

The *mlnx_add_kernel_support.sh* script can be executed directly from the *mlnxofedinstall* script. For further information, please see *--add-kernel-support* option.

Example

The following command will create a *MLNX_OFED_LINUX* ISO image for RedHat 6.3 under the */tmp* directory.

```
# ./MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64/mlnx_add_kernel_support.sh -m
/tmp/MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64/ --make-tgz
Note: This program will create MLNX_OFED_LINUX TGZ for rhel6.3 under /tmp directory.
All Mellanox, OEM, OFED, or Distribution IB packages will be removed.
Do you want to continue?[y/N]:y
See log file /tmp/mlnx_ofed_iso.21642.log
Building OFED RPMs. Please wait...
Removing OFED RPMs...
Created /tmp/MLNX_OFED_LINUX-x.x-x-rhel6.3-x86_64-ext.tgzhost1
```

The script adds the following lines to */etc/security/limits.conf* for the userspace components such as MPI:

– `* soft memlock unlimited`

– `* hard memlock unlimited`

These settings set the amount of memory that can be pinned by a user space application to unlimited. If desired, tune the value unlimited to a specific amount of RAM.

For your machine to be part of the InfiniBand/VPI fabric, a Subnet Manager must be running on one of the fabric nodes. At this point, Mellanox OFED for Linux has already installed the OpenSM Subnet Manager on your machine.

4.3.1.1 Installation procedure

Step 1

Login to the installation machine as *root*.

Step 2

Mount the ISO image on your machine.

```
host1# mount -o ro,loop MLNX_OFED_LINUX-<ver>-<OS label>-<CPU arch>.iso /mnt
```

Step 3

Run the installation script.

```
mnt/mnt/mlnxofedinstall --without-fw-update -force
Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs
This program will install the MLNX_OFED_LINUX package on your machine.
Note that all other Mellanox, OEM, OFED, or Distribution IB packages will be removed.
Uninstalling the previous version of MLNX_OFED_LINUX
Starting MLNX_OFED_LINUX-x.x.x installation ...
........
........
Installation finished successfully.
Attempting to perform Firmware update...
Querying Mellanox devices firmware ...
```

Step 4

Reboot the machine if the installation script performed firmware updates to your network adapter hardware. Otherwise, restart the driver by running: `/etc/init.d/openibd restart`

Step 5

(InfiniBand only) Run the *hca_self_test.ofed* utility to verify whether or not the InfiniBand link is up. The utility also checks for and displays additional information such as:

- HCA firmware version
- Kernel architecture
- Driver version
- Number of active HCA ports along with their states
- Node GUID

For more details on *hca_self_test.ofed*, see the file *hca_self_test.readmeunder docs/*.

After the installer completes, information about the Mellanox OFED installation such as prefix, kernel version, and installation parameters can be retrieved by running the *command /etc/infiniband/info*.

Most of the Mellanox OFED components can be configured or reconfigured after the installation by modifying the relevant configuration files. See the relevant chapters in this manual for details.

The list of the modules that will be loaded automatically upon boot can be found in the */etc/infiniband/openib.conf* file.

4.3.1.2 Installation result

Software

Most of *MLNX_OFED* packages are installed under the */usr* directory except for the following packages which are installed under the */opt* directory:

- *openshmem*, *bupc*, *fca* and *ibutils*

The kernel modules are installed under

- */lib/modules/`uname -r`/updates* on SLES
- */lib/modules/`uname -r`/extra/mlnx-ofa_kernel* on RHEL

Firmware

The firmware of existing network adapter devices will be updated if the following two conditions are fulfilled:

- The installation script is run in default mode; that is, without the option *--without-fw-update*

– The firmware version of the adapter device is older than the firmware version included with the Mellanox OFED ISO image

  If an adapter's Flash was originally programmed with an Expansion ROM image, the automatic firmware update will also burn an Expansion ROM image.

In case your machine has an unsupported network adapter device, no firmware update will occur and the error message below will be printed.

```
-I- Querying device ...
-E- Can't auto detect fw configuration file: ...
```

Please contact your hardware vendor for help on firmware updates.

4.3.1.3 Installation logging

While installing MLNX_OFED, the install log for each selected package will be saved in a separate log file.

The path to the directory containing the log files will be displayed after running the installation script in the following format:

`Logs dir: /tmp/MLNX_OFED_LINUX-`*version*`.`*PID*`.logs`

Example

```
Logs dir: /tmp/MLNX_OFED_LINUX-x.x-x.logs
```

4.3.1.4 Driver load upon system boot

Upon system boot, the Mellanox drivers will be loaded automatically.
To prevent automatic load of the Mellanox drivers upon system boot:

- Add the following lines to the */etc/modprobe.d/mlnx.conf* file.
  ```
  blacklist mlx4_core
  blacklist mlx4_en
  blacklist mlx5_core
  blacklist mlx5_ib
  ```
- Set `ONBOOT=no` in the */etc/infiniband/openib.conf* file.

4.3.1.5 mlnxofedinstall return codes

Table lists the mlnxofedinstall script return codes and their meanings.

Return	Code Meaning
0	The Installation ended successfully
1	The installation failed
2	No firmware was found for the adapter device
22	Invalid parameter
28	Not enough free space
171	Not applicable to this system configuration. This can occur when the required hardware is not present on the system.
172	Prerequisites are not met. For example, missing the required software installed or the hardware is not configured correctly.
173	Failed to start the mst driver

4.3.2 Port type management

ConnectX®-4 ports can be individually configured to work as InfiniBand ports. By default both ConnectX ports are initialized as InfiniBand ports. If you wish to change the port type, use the *mlxconfig* script after the driver is loaded.

For further information on how to set the port type in ConnectX®-4, please refer to the User Manual.

4.3.3 Uninstalling Mellanox OFED

Use the script */usr/sbin/ofed_uninstall.sh* to uninstall the Mellanox OFED package. The script is part of the ofed-scripts RPM.

4.4 Updating Firmware after Installation

If you wish to burn newer firmware, you have to download it from Fujitsu Technology Solutions Web site (*http://support.ts.fujitsu.com/download/* > *Drivers* & *Downloads*).
If you purchased the products from Fujitsu Japan, you download it from *http://jp.fujitsu.com/platform/server/primergy/downloads/*.

4.5 Subnet Manager

At least one Subnet Manager(SM) is present on each subnet. Each SM resides on a port of host channel adapter (HCA) or an InfiniBand Switch. When there are multiple SMs on a subnet, the master SM is one, and the remaining SMs becomes standby SMs. The master SM has a role to initialize and configure an InfiniBand subnet.

OpenSM is an InfiniBand compliant SM. The following sections describe how to use OpenSM.

4.5.1 Functions of OpenSM

OpenSM is one of Subnet Manager distributed with RedHat.

This chapter describes how to configure which servers becomes the master SM. Please set up the server which will be the master using the following reference. Regarding the other settings of OpenSM, please refer to the command help of *opensmd*.

Basically, a server with the lowest GUID will be the master SM within a subnet. if you intend to choose the master SM, you can configure the priority which range is from 0 to 15. 0 is the lowest and 15 is the highest.

Example for RHEL5:

```
# /etc/init.d/opensmd start -p 15
```

4.5.2 Confirmation that OpenSM is Running

Enter the following command to confirm that OpenSM is running:

```
# /etc/init.d/opensmd status
```

If OpenSM is running properly, the result of this command displays *running*.

Enter the following command to run OpenSM, if the result displays *stopped*:

```
# /etc/init.d/opensmd start
```

4.5.3 Confirmation of the master SM

Enter *ibstat* command on the system server, then "Base LID" and "SM LID" are displayed.

The IB HCA card on which these two LIDs are the same is the master SM.

5 Troubleshooting

5.1 General

Server unable to find the adapter	– Ensure that the adapter is placed correctly – Make sure the adapter slot and the adapter are compatible – Install the adapter in a different PCI Express slot – Use the drivers that came with the adapter or download the latest – Make sure your motherboard has the latest BIOS – Try to reboot the server
The adapter no longer works	– Reseat the adapter in its slot or a different slot, if necessary – Try using another cable – Reinstall the drivers for the network driver files may be damaged or deleted – Reboot the server
Adapters stopped working after installing another adapter	– Try removing and re-installing all adapters – Check that cables are connected properly – Make sure your motherboard has the latest BIOS
Link indicator light is off	– Ensure that adapter driver/s is loaded – Try another port on the switch – Make sure the cable is securely attached – Check your are using the proper cables that do not exceed the recommended lengths – Verify that your switch and adapter port are compatible
Link light is on, but with no communication established	– Check that the latest driver is loaded – Check that both the adapter and its link are set to the same speed and duplex settings

5.2 Linux

Environment Information	`cat/etc/issue`
	`uname -a`
	`cat/proc/cupinfo | grep` '*model name*' `| uniq`
	`ofed_info | head –1`
	`ifconfig -a`
	`ethtool` *<IB interface>*
	`ethtool -i` *<IB interface_of_Mellanox_port_num>*
	`ibdev2netdev`

Card Detection	`lspci | grep -i Mellanox`

Mellanox Firmware Tool (MFT)	Refer to the User Manual for installation instructions.
	Once installed, run:
	
	`mst start`
	`mst status`
	`flint -d` *<mst_device>* `q`

Ports Information	`ibstat` `lbv_devinfo`
Collect Log File	`/var/log/messages` `dmesg > system.log`

6 Appendix

6.1 A.1 EMC Certification Statements

	PIB EP 100Gb 1 Port EDR	PIB EP 100Gb 2 Port EDR
FCC	yes	yes
VCCI	yes	yes
EN	yes	yes
ICES	yes	yes
CE	yes	yes
CB	yes	yes
cTUVus	yes	yes
KCC	yes	yes
C-TICK	yes	yes
CCC	exemption	exemption
GOST-R	yes	yes

Table 4: Certification status

6.1.1 FCC Statements (USA)

Class A Statements:

- **§15.21**

  Statement

  CAUTION!

  Changes or modifications to this equipment not expressly approved by the party responsible for compliance could void the user's authority to operate the equipment.

- **§15.105(a)**

  Statement

  This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

6.1.2 EN Statements (Europe)

EN55022 Class A Statement:

CAUTION!

This is a class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

6.1.3 ICES Statements (Canada)

Class A Statement:

This Class A digital apparatus complies with Canadian ICES-003.

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

6.1.4 VCCI Statements (Japan)

Class A Statement:

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Translation - This is a Class A product based on the standard of the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). If this equipment is used in a domestic environment, radio interference may occur, in which case the user may be required to take corrective actions.

6.1.5 KCC Certification (Korea)

Korea's "Regulation for Certification of Information and Communication Equipment," requires EMC testing and certification for many electronic products. Korean EMC certifications are issued by Radio Research Laboratory (RRL), which is organized under the Ministry of Information and Communications (KCC). EMC testing includes electromagnetic emissions (EMI) and susceptibility (EMS). Certified equipment is labeled with the KCC mark and certification number.

이 기기는 업무용으로 전자파적합등록을 한 기기이오니 판매자 또는 사용자는 이 점을 주의하시기 바라며 만약 잘못 판매 또는 구입하였을 때에는 가정용으로 교환하시기 바랍니다.

Translation:
Class A Device
This device is registered for EMC requirements for industrial use. The seller or buyer should be aware of this. If this type was sold or purchased by mistake, it should be replaced with a residential-use type.

6.2 Specifications

PIB EP 100Gb 1 Port EDR

Physical Size:	2.71 in. x5.6 in. (68.90mm x 142.2 mm)	
	Connector:	QSFP28 InfiniBand (Copper and optical)
Protocol Support	InfiniBand:	IBTA v1.3
	Auto-Negotiation[1]:	1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane),QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port
	Data Rate:	Up to 100Gb/s EDR– InfiniBand
	PCI Express Gen3:	SERDES @ 8.0GT/s, 16 lanes
Power and Environmental	Voltage: 12V, 3.3V	
	Typ Power:	Passive Cables 14.25W
	Max Power:	Passive Cables 16.48W
	Max power available through QSFP28 port:	3.5W
	Temperature:	Operational 0°C to 55°C
		Non-operational -40°C to 70°C
	Humidity:	90% relative humidity[2]
	Air Flow[3]:	600LFM in case of 3.5W QSFP cable
		350LFM in case of 1.5W QSFP cable
Regulatory	Safety:	IEC/EN 60950-1:2006
		ETSI EN 300 019-2-2
		IEC 60068-2- 64, 29, 32
	RoHS:	RoHS-R6

1 The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another Mellanox InfiniBand product.

2 For both operational and non-operational states.

3 Air flow is measured ~1” from the heat sink between the heat sink and the cooling air inlet.

PIB EP 100Gb 2 Port EDR

Physical Size:	2.71 in. x5.6 in. (68.90mm x 142.2 mm)	
	Connector:	QSFP28 InfiniBand (Copper and optical)
Protocol Support	InfiniBand:	IBTA v1.3
	Auto-Negotiation[1]:	1X/2X/4X SDR (2.5Gb/s per lane), DDR (5Gb/s per lane),QDR (10Gb/s per lane), FDR10 (10.3125Gb/s per lane), FDR (14.0625Gb/s per lane), EDR (25Gb/s per lane) port
	Data Rate:	Up to 100Gb/s EDR– InfiniBand
	PCI Express Gen3:	SERDES @ 8.0GT/s, 16 lanes
Power and Environmental	Voltage: 12V, 3.3V	
	Typ Power:	Passive Cables 16.17W
	Max Power:	Passive Cables 18.43W
	Max power available through QSFP28 port:	3.5W
	Temperature:	Operational 0°C to 55°C
		Non-operational -40°C to 70°C
	Humidity:	90% relative humidity[2]
	Air Flow[3]:	600LFM in case of 3.5W QSFP cable
		350LFM in case of 1.5W QSFP cable
Regulatory	Safety:	IEC/EN 60950-1:2006
		ETSI EN 300 019-2-2
		IEC 60068-2- 64, 29, 32
	RoHS:	RoHS-R6

1 The ConnectX-4 adapters supplement the IBTA auto-negotiation specification to get better bit error rates and longer cable reaches. This supplemental feature only initiates when connected to another Mellanox InfiniBand product.

2 For both operational and non-operational states.

3 Air flow is measured ~1” from the heat sink between the heat sink and the cooling air inlet.